

\mathcal{L}_1 Adaptive Flight Control System: Flight Evaluation and Technology Transition

Enric Xargay*, Naira Hovakimyan†

University of Illinois at Urbana–Champaign, Urbana, IL 61801

Vladimir Dobrokhodov‡, Isaac Kaminer§

Naval Postgraduate School, Monterey, CA 93943

Irene M. Gregory¶

Chengyu Cao ||

NASA Langley Research Center, Hampton, VA 23681

University of Connecticut, Storrs, CT 06269

Certification of adaptive control technologies for both manned and unmanned aircraft represent a major challenge for current Verification and Validation techniques. A (missing) key step towards flight certification of adaptive flight control systems is the definition and development of analysis tools and methods to support Verification and Validation for nonlinear systems, similar to the procedures currently used for linear systems. In this paper, we describe and demonstrate the advantages of \mathcal{L}_1 adaptive control architectures for closing some of the gaps in certification of adaptive flight control systems, which may facilitate the transition of adaptive control into military and commercial aerospace applications. As illustrative examples, we present the results of a piloted–simulation evaluation on the NASA AirSTAR flight test vehicle, and results of an extensive flight test program conducted by the Naval Postgraduate School to demonstrate the advantages of \mathcal{L}_1 adaptive control as a verifiable robust adaptive flight control system.

I. Introduction

Inner–loop adaptive flight control systems (FCSs) may provide the opportunity to improve aircraft performance and reduce pilot’s workload at challenging flight envelope conditions or in the event of control surface failures and vehicle damage. However, implementing adaptive control technologies can increase the complexity of FCSs beyond the capability of current Verification and Validation (V&V) processes. This fact, combined with the criticality of inner–loop flight control systems, renders difficult the transition of these technologies into military and commercial applications. A good overview on verification, validation, and certification challenges for the transition of adaptive control to safety–critical aerospace applications can be found in [1,2]. Programs like NASA’s “Integrated Resilient Aircraft Control” (IRAC) and Wright–Patterson AFRL’s “Certification Techniques for Advanced Flight Critical Systems” (CerTA FCS) represent an effort to advance the state–of–the–art in adaptive control algorithms, to analyze the deficiencies of current V&V practices, and to advance airworthiness certification of adaptive flight control systems. In particular, these two programs have significantly contributed to the ongoing efforts in the development, flight verification and validation, and transition of \mathcal{L}_1 adaptive control from a theoretical research field into a viable and reliable technology towards improving the robustness and performance of advanced flight control systems.

It is well known that conventional theory of adaptive control has limited analysis framework for its robustness and performance guarantees. The key limitation of the adaptive control architectures developed

*Graduate Student, Dept. of Aerospace Engineering, Student Member AIAA; xargay@illinois.edu.

†Professor, Dept. of Mechanical Science and Engineering, Associate Fellow AIAA; nhovakim@illinois.edu.

‡Research Assistant Professor, Dept. of Mech. & Astronautical Eng., Member AIAA; vldobr@nps.edu.

§Professor, Dept. of Mech. & Astronautical Eng., Member AIAA; kaminer@nps.edu.

¶Senior Aerospace Research Engineer, Dynamic Systems and Controls Branch, MS 308, Associate Fellow AIAA; irene.m.gregory@nasa.gov.

||Assistant Professor, Dept. of Mechanical Engineering, Member AIAA; ccao@engr.uconn.edu.

during the last forty years is the lack of predictability in the interaction of the adaptation process with the performance and robustness characteristics of the closed-loop adaptive system. As a result, the standard well-known trade-off between performance and robustness in feedback control system design becomes a “three-party” trade-off between adaptation, performance, and robustness.³ In general, increasing the speed of adaptation leads to loss of robustness, whereas slow adaptation leads to lack of transient characterization with respect to changes in reference inputs, initial conditions, uncertainty, etc. Due to this inherent limitation of conventional adaptive control architectures and despite the vast improvements in adaptive control design methods observed during the years, adaptive control has largely remained a tool for adapting only to slowly varying uncertainties. Moreover, the lack of guidelines to solve the trade-off between adaptation, performance, and robustness makes the design of conventional adaptive controllers an overly challenging problem, which is being commonly resolved by either computationally expensive Monte-Carlo simulations or trial-and-error methods following some empirical guidelines or engineering intuition. This, however, leads to long development time and high certification costs.

In order to overcome some of these difficulties and improve transient performance with satisfactory robustness margins, some adaptive control architectures rely on *persistence of excitation* assumptions or on the injection of *persistently exciting signals* on top of the reference commands or feedback signals. In inner-loop FCS design, however, high persistent excitation can potentially cause undesirable aircraft behavior, induce significant cross-coupling between different axes, make adverse airplane-pilot coupling more likely, or excite unmodeled dynamics such as flexible modes. Therefore, trying to enforce persistence of excitation could actually hurt the robustness of adaptive FCSs and, in the case of impaired aircraft, even aggravate the situation further. In this sense, even if significant efforts are underway to analyze pilot interactions with adaptive control systems, persistence of excitation seems to be a no-go decision for development of reliable FCSs. The lack of analytical quantification of the relationship between the rate of adaptation, the transient response, and the robustness margins has also resulted in the development of *gain-scheduled* designs of adaptive controllers, examples of which are the flight tests of the late 90’s by the Air Force and Boeing.^{4,5} The design of such controllers heavily relies on Monte-Carlo simulations and, hence, the approach seems to defeat the purpose of adaptation, which appeared as a potential alternative to gain-scheduling in flight control system design.

In contrast to these proposed adaptive control architectures, \mathcal{L}_1 adaptive control overcomes the difficulties encountered in the development of adaptive control algorithms, provides a proper framework for V&V, and closes some of the gaps in certification of adaptive flight control systems. This may facilitate the transition of adaptive control into aerospace applications. The main benefit of \mathcal{L}_1 adaptive control is its *fast* and *robust* adaptation, which does not interact with the trade-off between performance and robustness.⁶ In fact, the architectures of \mathcal{L}_1 adaptive control theory have *guaranteed transient performance* and *guaranteed robustness* in the presence of *fast adaptation*, without introducing or enforcing persistence of excitation, without any gain scheduling in the controller parameters, and without resorting to high-gain feedback. In these architectures, the speed of adaptation is limited only by the available CPU (hardware), while robustness can be resolved via conventional methods from classical and robust control. To make things more precise, the philosophy of \mathcal{L}_1 adaptive control architectures is to obtain the estimate of the uncertainties via a fast estimation scheme and define the control signal as the output of a low-pass linear filter (commonly denoted $C(s)$ in \mathcal{L}_1 architectures), which not only guarantees that the control signal stays in the low-frequency range in the presence of fast adaptation and large reference inputs, but also leads to separation between adaptation and robustness, and moreover it defines the trade-off between performance and robustness. In particular, proper filter design leads to desired transient performance for the system’s input and output signals simultaneously, in addition to steady-state tracking.⁷ It has also guaranteed, analytically provable, bounded-away-from-zero time-delay margin.⁸ Moreover, \mathcal{L}_1 adaptive control theory’s systematic design procedures significantly reduce the tuning effort required to achieve desired closed-loop performance, particularly while operating in the presence of uncertainties and failures. It is also important to mention that the fast adaptation capability of \mathcal{L}_1 adaptive control architectures allows for control of time-varying nonlinear systems by adapting two parameters only,⁹ which is in contrast to the use of on-line neural networks which may have problems of computational predictability^{1,10} and for which the only V&V techniques are based on brute-force Monte-Carlo testing.

In brief, the \mathcal{L}_1 adaptive control theory shifts the tuning issue from the selection of the adaptive gain of the underlying *nonlinear* gradient minimization algorithm to determining the structure/bandwidth for a *linear* low-pass filter in the control channel. The main features of the \mathcal{L}_1 adaptive control theory, proven in

theory and consistently verified in experiments, can thus be summarized as follows:

- Guaranteed robustness in the presence of fast adaptation;
- Separation (decoupling) between adaptation and robustness;
- Guaranteed transient response, *without* resorting to
persistence of excitation type assumptions,
high-gain feedback,
gain-scheduling of the controller parameters,
control reconfiguration;
- Guaranteed (bounded-away-from-zero) time-delay margin;
- Uniform *scaled transient response* dependent on (admissible) changes in
initial conditions,
system uncertainty,
and reference inputs.

With these features, the \mathcal{L}_1 adaptive control architectures reduce the performance limitations to hardware limitations and provide a suitable framework for development of theoretically justified tools for V&V of feedback systems, and thus supports flight control design in the presence of damage, failures, and other adverse conditions.

In this paper, we present the results of our ongoing efforts in the development, flight verification and validation, and technology transition of \mathcal{L}_1 adaptive control to aerospace applications. In particular we present (i) preliminary results of the application and flight evaluation of an \mathcal{L}_1 adaptive flight control system to the Airborne Subscale Transport Aircraft Research (AirSTAR) system at the NASA Langley Research Center; and (ii) the development of a unified framework for design, implementation, verification and validation of \mathcal{L}_1 adaptive flight critical control systems conducted by the Naval Postgraduate School (NPS). The purpose of this work is to demonstrate the advantages of \mathcal{L}_1 adaptive control as a *verifiable* robust adaptive control architecture with the potentiality of reducing flight control design costs and facilitating the transition of adaptive control into advanced flight control systems.

The paper is organized as follows. Section II presents preliminary results on the development and application of \mathcal{L}_1 adaptive control architectures to the design of inner-loop flight control systems. First, we present preliminary results of a piloted-simulation evaluation on the AirSTAR flight test vehicle high-fidelity simulator. We provide specific details about the flight control system design, and a detailed description of the results. And second, we present hardware-in-the-loop simulations and flight test results by NPS, in which we analyze the robustness and performance of an \mathcal{L}_1 adaptive output feedback controller in the presence of unmodeled dynamics. In particular, we demonstrate the performance of the closed-loop adaptive system in nominal conditions by varying frequency of the reference signal. We also illustrate the robustness of the adaptive algorithm in the presence of high-frequency bending modes. Finally, Section III summarizes the key results and presents the main conclusions.

II. \mathcal{L}_1 Adaptive Control in Transition

A. \mathcal{L}_1 Adaptive Control Design for NASA AirSTAR Flight Test Vehicle

The research control law developed for the NASA's Generic Transport Model (GTM) aircraft has as its primary objective achieving tracking for a variety of tasks with guaranteed stability and robustness in the presence of uncertain dynamics, such as changes due to rapidly varying flight conditions during standard maneuvers, and unexpected failures. All of this must be achieved while providing Level I handling qualities under nominal as well as adverse flight conditions. In particular, one essential objective for safe flight under adverse conditions is for the aircraft never to leave the extended α - β flight envelope, for which accurate models of the aircraft are known (see Figure 1). Outside the boundary of this extended flight envelope, the aerodynamic models available are usually obtained by *extrapolation* of wind-tunnel test data, implying

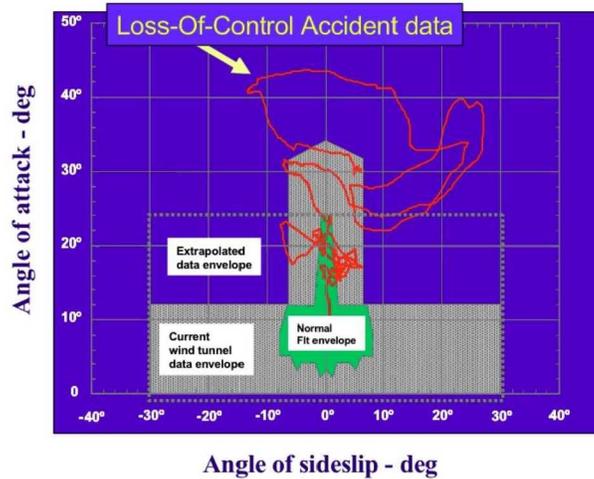


Figure 1: Loss of control accident data relative to angle of attack and angle of sideslip (from [11])

that these models are highly uncertain. This fact suggests that pilots might not be correctly trained to fly the aircraft in these regimes (or, in the case of unmanned vehicles, the guidance loops might not be properly designed for safe recovery). Moreover, it does not seem reasonable to rely on a flight control system to compensate for the uncertainties in these flight conditions, as aircraft dynamic controllability is not even guaranteed in such regimes. Consequently, the adaptive controller should learn fast enough to keep the aircraft within this extended flight envelope. This implies that the control law action in the initial 2–3 seconds after initiation of an adverse condition is the key to safe flight.

The \mathcal{L}_1 control system developed for this application is a three axes angle of attack (α), roll rate (p)–sideslip angle (β) command augmentation system. Specifically, we use a theoretical extension of the \mathcal{L}_1 adaptive control theory that compensates for both *matched* as well as *unmatched* dynamic uncertainties.¹² For inner-loop FCS design, the effects of slow outer-loop variables (e.g. airspeed, pitch angle, bank angle) may appear as unmatched uncertainties in the dynamics of the fast inner-loop variables we are trying to regulate (e.g. angle of attack, sideslip angle, roll rate). Also, unmodeled nonlinearities, cross-coupling between channels, and dynamic asymmetries may introduce unmatched uncertainties in the inner-loop system dynamics. If the design of the inner-loop FCS does not account for these uncertainties, their effect in the inner-loop dynamics will require continuous compensation by the pilot, thereby increasing the pilot’s workload. Therefore, automatic compensation for the undesirable effects of these unmatched uncertainties on the output of the system is critical to achieve desired performance, reduce pilot’s workload, and improve the aircraft’s handling qualities.

It is important to note also that the \mathcal{L}_1 adaptive flight control system provides a systematic framework for adaptive controller design that allows for explicit enforcement of MIL-Standard requirements and significantly reduces the tuning effort required to achieve desired closed-loop performance. This in turn reduces the design cycle time and the development costs. In particular, the design of the \mathcal{L}_1 adaptive flight control system for the GTM is based on the linearized dynamics of the aircraft at an (equivalent) airspeed of 80 knots and at an altitude of 1000 ft. Since the airplane is Level I at this flight condition, the nominal desired dynamics of the (linear) state-predictor were chosen to be similar to those of the actual airplane; only some additional damping was added to both longitudinal and directional dynamics, while the lateral dynamics were set to be slightly faster than the original ones in order to satisfy performance specifications. The state-predictor was scheduled to specify different performance requirements at special flight regimes (high-speed regimes and high- α regimes). In order to improve the handling qualities of the airplane, a linear prefilter was added to the adaptive FCS so as to ensure desired decoupling properties as well as desired command tracking performance. Overdamped second-order low-pass filters with unity DC gain were used in all the control channels, while their bandwidths were set to ensure (at least) a time-delay margin of 100 ms and a gain margin of 6 dB. Finally, the adaptation sampling time was set to $\frac{1}{600}$ s, which corresponds to the maximum integration step allowed in the AirSTAR flight control computer. We notice that the same

control parameters for the prefilter, the low-pass filters, and the adaptation sampling time were used across the entire flight envelope with no scheduling or reconfiguration.

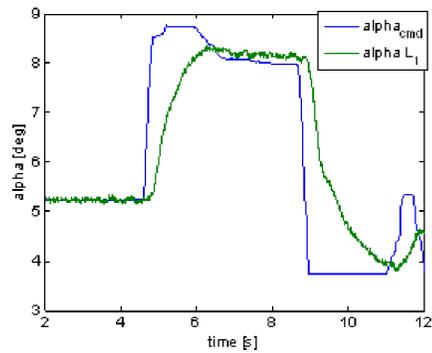
Next we present preliminary results of a piloted-simulation evaluation on the GTM aircraft high-fidelity simulator, which includes full nonlinear, asymmetric aerodynamics, actuator dynamics, sensor dynamics including noise, biases, and scaling factors, and other nonlinear elements typical of these kind of simulators. This piloted-simulation evaluation was part of development and profile planning of flight test tasks. The tasks were flown with no training or repeatability and thus are considered a preliminary evaluation. The GTM aircraft has three flight control modes. *Mode 1* is the revisionary stick to surface control, the aircraft is Level I in this configuration; *Mode 2* is referred to as the *baseline control law* and is an α -command, p - β stability augmentation system; *Mode 3* is the *research control law* which is the \mathcal{L}_1 adaptive α , p - β command control. The baseline α -command is an LQR PI based controller and is designed to the operational limit of $\alpha \leq 10$ deg. It currently serves as the operational baseline feedback control law for the GTM. Hence, comparison between the \mathcal{L}_1 and the baseline α -command can be only performed for $\alpha \leq 10$ deg, comparisons for higher α range are performed with the stick to surface control law (Mode 1).

ANGLE OF ATTACK CAPTURES IN THE PRESENCE OF STATIC STABILITY DEGRADATION. The task is trim at $V = 80$ KEAS and $Alt = 1200$ ft and then capture $\alpha = 8$ deg within 1 second, hold for 2 seconds, with α -desired ± 1 deg, α -adequate ± 2 deg. This task was repeated for various levels of static stability expressed as $\Delta C_{m\alpha}$ and ranging from 0 to 100%, i.e. from nominal stable aircraft to neutral static longitudinal stability. The change in $\Delta C_{m\alpha}$ is achieved by using both inboard elevator sections, scheduled with α , to produce a destabilizing effect. These two elevator sections also become unavailable to the control law. In a sense, it is a double fault—a destabilized aircraft and a reduction in control power in the affected axis.

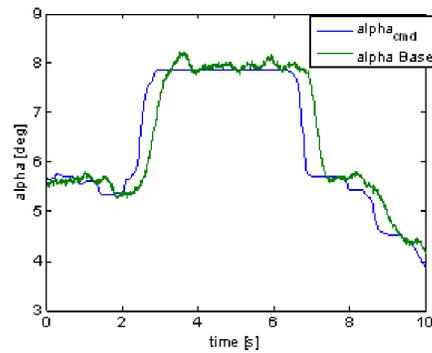
This longitudinal task was evaluated for the \mathcal{L}_1 adaptive control law and the baseline, both of which being α -command response type. The performance of both control laws is provided in Figure 2. For nominal GTM aircraft, performance of both control laws is very similar, as illustrated in Figures 2a–2b, and a solid Level I flying qualities (FQ) according to pilot comments. However, as the static stability is decreased by 50% (Figures 2c–2d), the performance of the baseline controller degrades to high Level II (CHR 4), while the \mathcal{L}_1 adaptive controller remains solid Level I FQ. With $\Delta C_{m\alpha} = 75\%$ (Figures 2e–2f), the \mathcal{L}_1 adaptive controller remains predictable and Level I, while the baseline performance degrades to achieving only adequate performance (Level II, CHR 5). At the point of neutral static stability (Figures 2g–2h), the \mathcal{L}_1 adaptive controller is still described as predictable, but does experience some oscillations, and its performance is reduced to high Level II (CHR 4), while the baseline controller is described as PIO prone and FQ are reduced to Level 3 bordering on uncontrollable CHR 10. One point we would like to emphasize is that the performance of the \mathcal{L}_1 adaptive control law was found predictable by the pilot for all levels of static stability.

HIGH ANGLE OF ATTACK CAPTURES. One of the several research objectives for the AirSTAR facility is to identify high angle of attack dynamics and verify these against obtained wind tunnel and CFD data. In order to do so, the GTM aircraft must be able to safely fly at the very edges of the attainable flight envelope. Part of the scheduled flight test is the high α envelope expansion. In addition, the GTM exhibits highly nonlinear *pitch break* phenomena for $12 \leq \alpha \leq 18$ deg. In other words, if in open loop the aircraft reaches $\alpha = 12$ deg, it will pitch up and could be recovered only once it reaches $\alpha = 18$ deg. Thus the envelope expansion is flown in the revisionary stick-to-surface control mode with 2 deg- α increments from $\alpha = 18$ deg to $\alpha = 28$ deg. On the other hand, the \mathcal{L}_1 adaptive controller is expected to perform α capture task for entire post-stall region starting at $\alpha = 12$ deg.

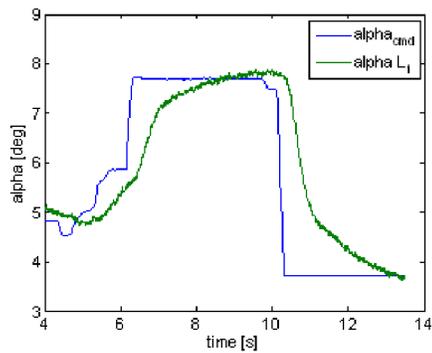
The aerodynamics in the post-stall region are nonlinear and increasingly asymmetric with increased α . For $12 \leq \alpha \leq 20$ deg the aerodynamics are expected to be asymmetric, the roll damping (C_{lp}) is expected to be low, and nose-slice roll-off is also expected. Beyond $\alpha = 28$ deg, in addition to aerodynamic asymmetry and low roll damping, a pronounced nose-slice due to $C_{n\beta}$ is expected. The α capture performance for the \mathcal{L}_1 adaptive controller is illustrated in Figure 3. The task is, starting in trim at $V = 80$ KEAS, to capture the indicated α at the rate of $3 \frac{\text{deg}}{\text{s}}$, hold for 4 seconds, with α -desired ± 1 deg, α -adequate ± 2 deg. Note that the desired and the adequate criteria are the same as for the α capture in the linear region; additionally, holding for 4 seconds would expose any control law to instability in this “pitch break” region. Due to the nature of expected dynamic behavior in the high α region, β and ϕ are additional variables of interest plotted in Figure 3. The \mathcal{L}_1 adaptive controller performance was judged as close to nominal α capture task. For $\alpha = 20$ deg case (Figure 3b), the approach was at high pitch rate with V_{\min} getting into the 30’s (KEAS);



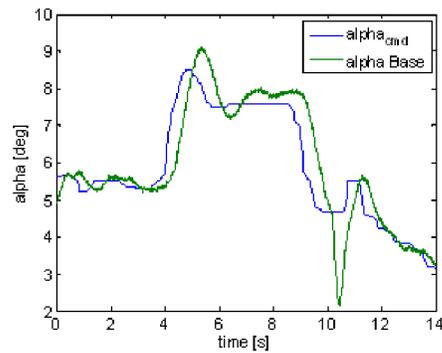
(a) $\Delta C_{m\alpha} = 0\%$ - \mathcal{L}_1 Adaptive



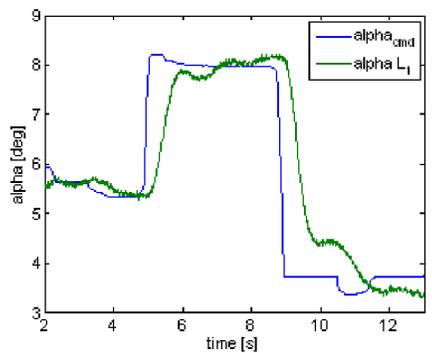
(b) $\Delta C_{m\alpha} = 0\%$ - Baseline



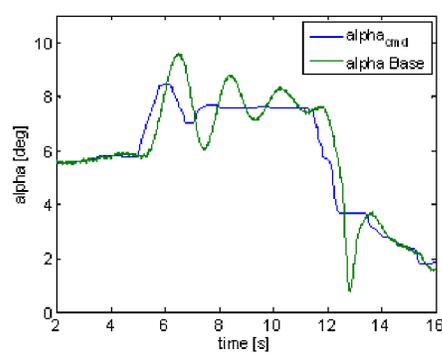
(c) $\Delta C_{m\alpha} = 50\%$ - \mathcal{L}_1 Adaptive



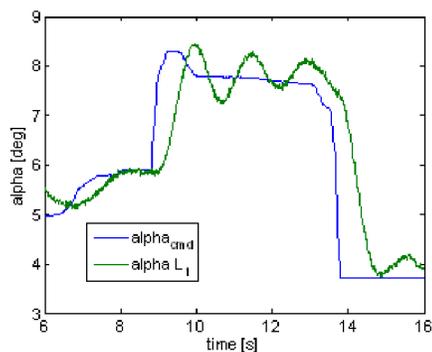
(d) $\Delta C_{m\alpha} = 50\%$ - Baseline



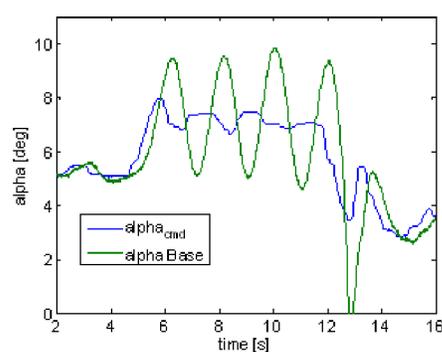
(e) $\Delta C_{m\alpha} = 75\%$ - \mathcal{L}_1 Adaptive



(f) $\Delta C_{m\alpha} = 75\%$ - Baseline



(g) $\Delta C_{m\alpha} = 100\%$ - \mathcal{L}_1 Adaptive



(h) $\Delta C_{m\alpha} = 100\%$ - Baseline

Figure 2: Angle of attack capture task with variable static stability.

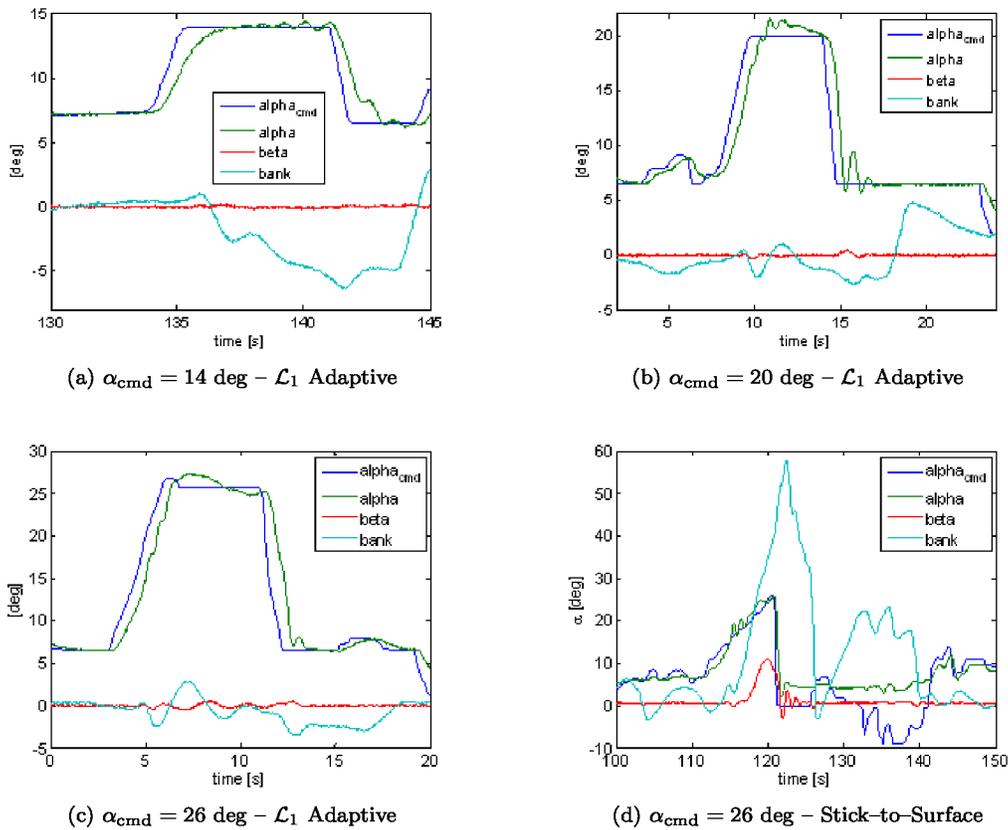


Figure 3: High α capture task.

this created some nascent small oscillations as α approached 20 deg. The $\alpha = 26$ deg case was performed less rapidly and the slight oscillations are no longer present (Figure 3c). For comparison purposes, the same α case for the revisionary stick-to-surface mode is shown in Figure 3d. Note the oscillatory nature of α as it follows the ramping α_{cmd} , also note the bank and sideslip angle excursions, which illustrate the expected roll-off (ϕ) and nose-slice (β) dynamics. Figure 4 provides another way of looking at the high α capture task by illustrating the coupling between these variables. Ideally, an α excursion would be completely decoupled from β and would produce a straight vertical line on the plot. The \mathcal{L}_1 adaptive controller produces $|\beta| \leq 1$ deg excursions, as shown on Figure 4a. On the other hand, the revisionary stick-to-surface control law shows significantly greater degree of coupling between the α - β axis, as illustrated in Figure 4b.

In Figure 4, one can see that the \mathcal{L}_1 adaptive controller keeps the airplane inside the *normal flight envelope* during all of the high- α maneuvers, thus ensuring controllability of the airplane during the whole task. The same cannot be said about the revisionary stick-to-surface control law, for which the airplane experiences high- α /high- β excursions.

SUDDEN ASYMMETRIC ENGINE FAILURE. This maneuver was performed unrehearsed once for each control law —reversion stick-to-surface mode and \mathcal{L}_1 adaptive controller. The results of this maneuver are used for qualitative comparison between the two control law responses with pilot in the loop and no training. The task starts the airplane climbing at ≈ 30 deg attitude and throttles at full power, then at some time the left throttle is reduced from 100% to 0% thrust in less than 0.5 s as illustrated in Figure 5. The left throttle transient occurs at approximately 12 s in Figure 5a and approximately 38 s in Figure 5b. The result is primarily a lateral-directional task since the sudden change in thrust induces rolling moment affecting p , ϕ and side force affecting β . These variables are co-plotted with throttle activity in Figures 5a and 5b. In particular, Figure 5a shows that the pilot is not able to stabilize the aircraft without assist. From Figure 5b, instead, it is evident that this sudden asymmetric thrust is a nonevent for the \mathcal{L}_1 adaptive controller especially from the stability perspective. The loss of control for the revisionary stick to surface control law is evident

from Figure 5c, which reveals α - β excursions outside the *extended flight envelope*. When the \mathcal{L}_1 flight control system is assisting the pilot, the airplane remains inside the *normal flight envelope* during the entire task, ensuring its controllability.

We would like to emphasize that the \mathcal{L}_1 adaptive controller was *not* redesigned or retuned in all of these scenarios, and that a *single set* of control parameters was used for all the piloted tasks and throughout the whole flight envelope. As stated earlier, only the reference model (state-predictor) is scheduled in order to specify different performance requirements at different flight regimes.

B. Rohrs' Counterexample in Flight

The lack of robustness exhibited by conventional model reference adaptive control (MRAC) was first pointed out in 1979 by Egardt¹³ and analyzed in greater detail by Rohrs et al, see [14]. In particular, Rohrs and coauthors constructed a counterexample where a first-order stable plant with two highly-damped unmodeled poles went unstable when driven by a reference sinusoid at the phase crossover frequency. In this section, we extend the setup introduced in [14] to flight test environment, i.e we replace the first-order nominal plant used there with a small unmanned aerial vehicle (UAV) controlled by a commercial autopilot (AP). The issues raised in [14] coincide with the objectives of IRAC-1.5 and, therefore, form an ideal outline to follow. Hence, an \mathcal{L}_1 adaptive output feedback controller is used to augment the nominal plant and is asked to track a sinusoidal signal of a specific frequency in the presence of modeling uncertainties.

The flight test setup used in this paper was first reported in [15], where the authors provided preliminary results on flight testing of \mathcal{L}_1 adaptive control system in the presence of UAV control surface failures. The control architecture included a commercial autopilot augmented by the output feedback \mathcal{L}_1 controller. The results obtained demonstrated that the \mathcal{L}_1 augmented system provides fast recovery to the sudden locked-in-place failures in either one of the ailerons or in the rudder, while the unaugmented system goes unstable. In the present paper, the results for the nominal plant consisting of a UAV and a commercial AP (see Figure 6) are extended to include “unmodeled dynamics” “injected” at the output of the plant, similar to the approach outlined in [14]. Two cases of unmodeled dynamics are considered. In the first case, these unmodeled dynamics are represented by a second-order transfer function that introduces significant uncertainty both at high and low frequencies, see Figure 7a. In the second case, uncertainty is given by a very lightly damped second-order system that represents, for example, a flexible body mode of an airplane, see Figure 7b (the natural frequency of the uncertainty is selected to be outside of the nominal system bandwidth). This last setup, in turn, allows us to validate if a small disturbance at the resonant frequency of the transfer function can propagate to the control signal and excite the bending mode.

1. Hardware-in-the-Loop Simulation Results

SECOND-ORDER UNMODELED DYNAMICS. In this section, we introduce artificially at the output of the plant (closed-loop UAV and AP) a second-order system representing unmodeled dynamics, similar to Rohrs' counterexample. In his counterexample, Rohrs considered the case of very well damped unmodeled dynamics at high frequencies to show that, even in the presence of apparently “harmless” uncertainties, the stability of adaptive controllers was not guaranteed. In this paper, nevertheless, we consider the case of more challenging uncertainties, with the objective of evaluating the performance of the \mathcal{L}_1 adaptive augmentation in the presence of unmodeled dynamics. In particular, we choose a low-damped second-order system with natural frequency equal to the bandwidth of the plant.

Next we present the results obtained for the same experiment with the \mathcal{L}_1 adaptive augmentation (Figure 8). As it can be seen, the closed-loop adaptive system is able to track precisely the output of the reference system to the reference signal at $\omega = 0.1 \frac{\text{rad}}{\text{s}}$. As the frequency of the reference signal increases and goes beyond the bandwidth of the low-pass filter $C(s)$, the tracking performance degrades as expected. The \mathcal{L}_1 adaptive controller keeps the DC gain of the closed-loop adaptive system close to 1 for the reference signals at frequencies $\omega = 0.1 \frac{\text{rad}}{\text{s}}$, $\omega = 0.5 \frac{\text{rad}}{\text{s}}$, and $\omega = 1 \frac{\text{rad}}{\text{s}}$, while the commands sent to the AP remain always bounded and without high-frequency content. Should one want to reduce the amplitude of the command from the adaptive controller to the AP, it would be enough to shrink the projection bounds for the adaptive law. This would limit as well the capabilities of the \mathcal{L}_1 adaptive controller. Also, one can observe that the response of the closed-loop adaptive system for every reference signal is very consistent and describes a well-defined and “clear” shape in the Lissajous plot.

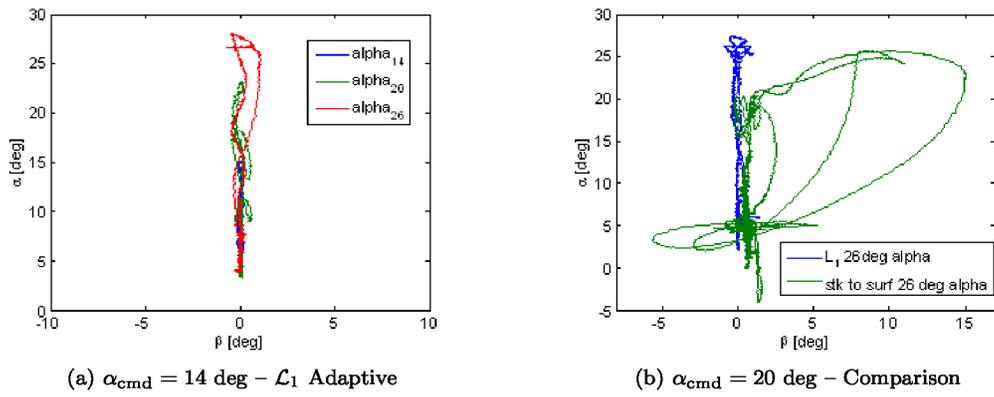


Figure 4: α - β excursions for high α capture task.

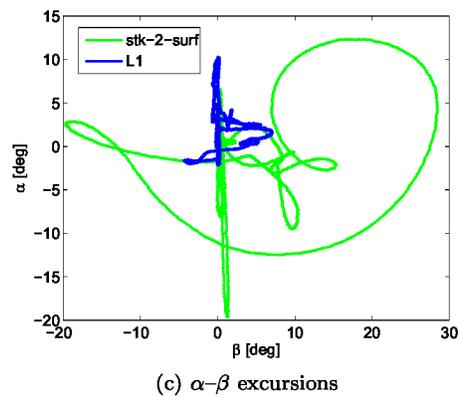
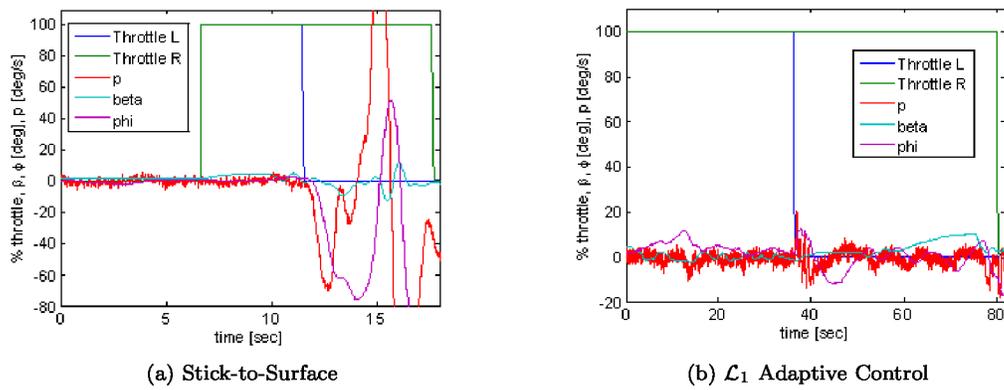


Figure 5: Asymmetric engine failure.

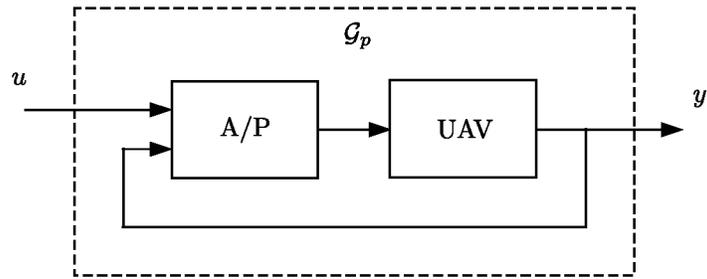
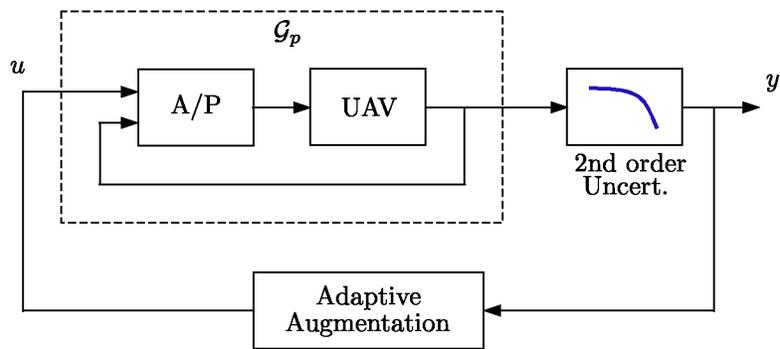
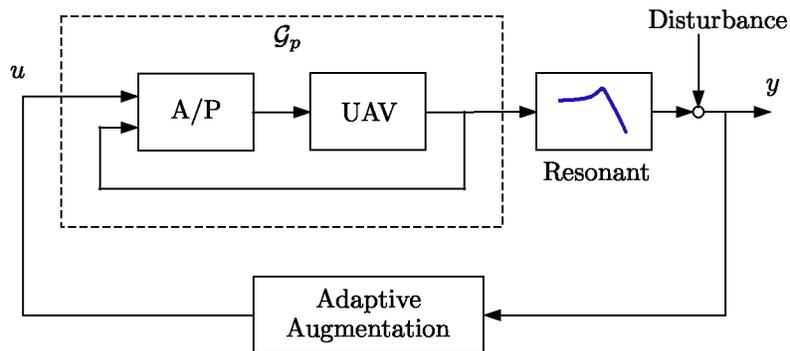


Figure 6: Nominal plant consisting of a UAV and a commercial AP.

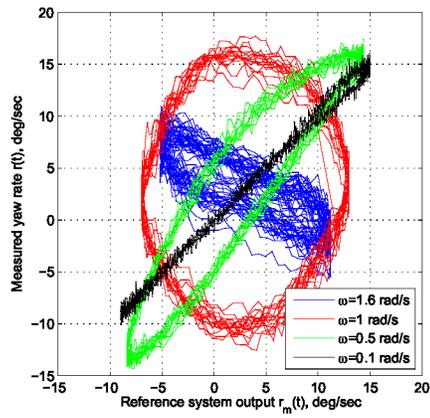


(a) Second-order transfer function at phase crossover frequency.

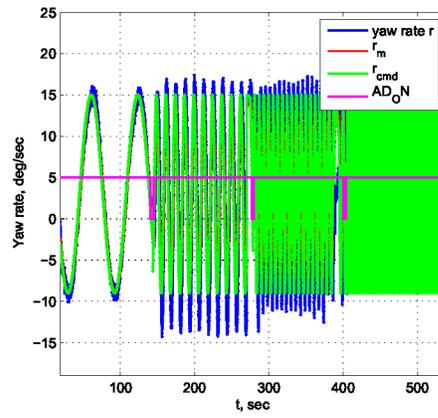


(b) High-Frequency Bending Mode

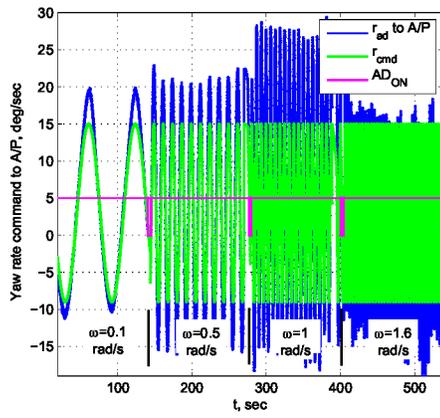
Figure 7: System with two cases of unmodeled dynamics and output disturbance



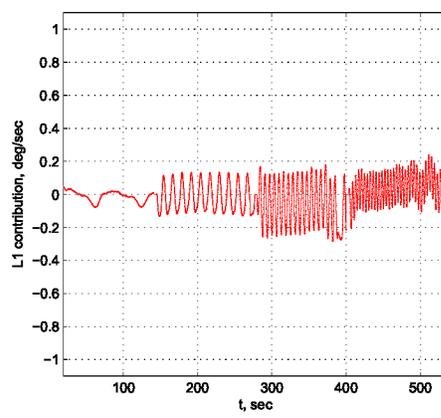
(a) Lissajous $r(r_m)$



(b) Tracking performance



(c) Adaptive command to AP



(d) \mathcal{L}_1 contribution

Figure 8: \mathcal{L}_1 . Closed-loop response in the presence of second-order unmodeled dynamics to biased sinusoidal reference signals.

BENDING MODE. In this section we extend the philosophy of the previous experiment to the case of high-frequency unmodeled dynamics with very low damping. This experiment has as a goal to reproduce the effects of bending modes in an airplane, and investigate the possible interaction between adaptive control and these modes. To this end, we choose a very low damped second-order transfer function with natural frequency equal to approximately 7 times the bandwidth of the plant.

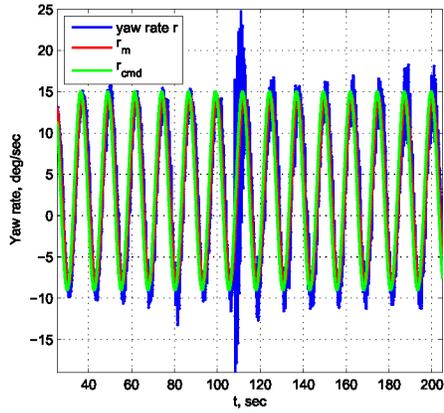
Next we present the results obtained with the \mathcal{L}_1 adaptive controller. At $t = 50$ s the bending mode is “injected”, which results in small amplitude oscillations at the output of the system due to the noise (see Figure 9). It can be seen that the contribution of the \mathcal{L}_1 adaptive controller stays in the low-frequency range. The disturbance is “injected” at $t = 108$ s. One can see that the \mathcal{L}_1 adaptive controller generates small oscillations in the control channel, which causes a 5 s transient at the output of the system with high-frequency oscillations. These oscillations in the contribution of the \mathcal{L}_1 controller are due to the fact that the projection bounds in the adaptive law are too tight to handle the initial transient produced by the disturbance. Figure 10 shows that the uncertainty estimate $\hat{\sigma}(t)$ hits the bounds of projection, and as a consequence, the error between the output of the state-predictor $\hat{r}(t)$ is not able to track the output of the plant $r(t)$. These oscillations could be easily avoided by increasing the bounds of the projection operator and, if necessary, increasing the adaptation gain to reduce the tracking error between the output of the state-predictor and the output of the plant. In the experiment, after these 5 s of transient, the oscillations damp down and the \mathcal{L}_1 adaptive controller recovers its initial performance. Although the \mathcal{L}_1 adaptive controller senses this high-frequency disturbance through the feedback signal, fast adaptation allows for rapid estimation of the uncertainties, $\hat{\sigma}(t)$. This estimation goes to the low-pass filter in the control law, which effectively attenuates this high-frequency signal, and therefore the \mathcal{L}_1 contribution remains within the low-frequency range. This results in the fact that the \mathcal{L}_1 controller is not “fighting” the AP. This experiment shows that, if properly tuned, the \mathcal{L}_1 adaptive controller guarantees that its contribution remains in the low-frequency range and preserves the attenuating properties of the nominal inner-loop, and thus the bending mode is not excited.

2. Flight Test Results

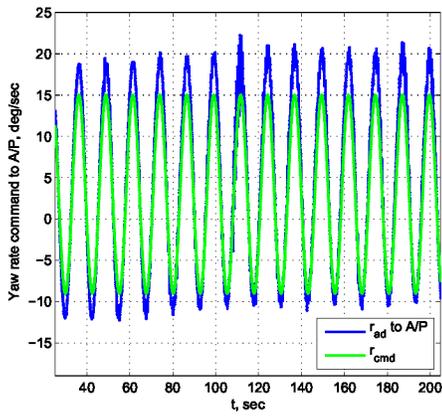
SECOND-ORDER UNMODELED DYNAMICS. The same second-order transfer function used in hardware-in-the-loop experiments was “injected” at the output of the actual UAV. Figure 11 presents the results obtained for the \mathcal{L}_1 adaptive augmentation. The system maintains stability during the whole flight and the control signal remains inside reasonable bounds during the experiment. As one would expect, since the frequency of the reference signal is well beyond the bandwidth of the low-pass filter $C(s)$ in the control law, the \mathcal{L}_1 adaptive controller is not able to recover desired performance of the closed-loop adaptive system. The response with \mathcal{L}_1 adaptive controller is consistent during the entire flight, and does not exhibit undesirable characteristics like bursting. In order to illustrate the degradation of performance as the frequency of the reference signal increases beyond the bandwidth of the low-pass filter, the closed-loop system with the second-order unmodeled dynamics and the \mathcal{L}_1 controller implemented onboard was driven with a set of biased sinusoidal reference signals at different frequencies. Figure 12 shows the results of these experiments. It can be seen that the output of the closed-loop adaptive system is able to track the output of the reference system for reference signals at low frequencies ($\omega = 0.3 \frac{\text{rad}}{\text{s}}$), and, as the frequency of the reference signal increases, the ideal Lissajous curve with slope-1 “inflates” progressively. This graceful degradation in the performance of the system is consistent with the theory of fast and robust adaptation, and results in a predictable response of the closed-loop adaptive system.

III. Conclusions

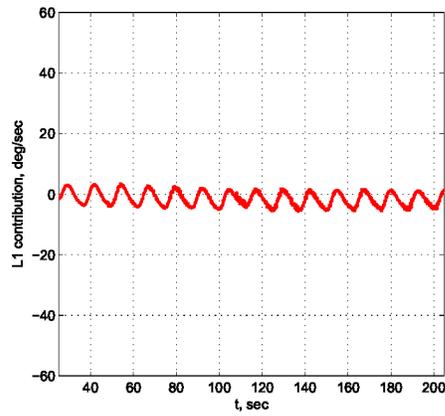
In this paper, we presented the benefits of \mathcal{L}_1 adaptive control for the development, flight verification and validation, and certification of adaptive flight control systems. In contrast to conventional adaptive control architectures, the recently developed \mathcal{L}_1 adaptive control provides a systematic framework for the design of adaptive controllers with performance and robustness guarantees. Moreover, the systematic design guidelines of \mathcal{L}_1 adaptive control may reduce the development cycle time and the design costs of adaptive flight control systems. Results from a piloted-simulation evaluation on the NASA AirSTAR vehicle, and hardware-in-the-loop simulations and flight tests conducted by NPS illustrate the performance and robustness characteristics of \mathcal{L}_1 adaptive flight control systems.



(a) Tracking performance

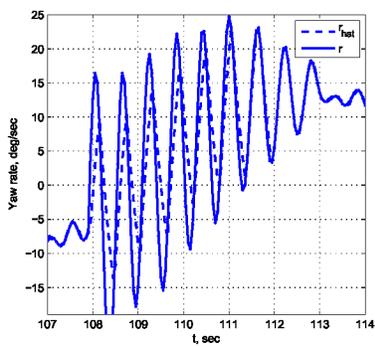


(b) Adaptive command to AP

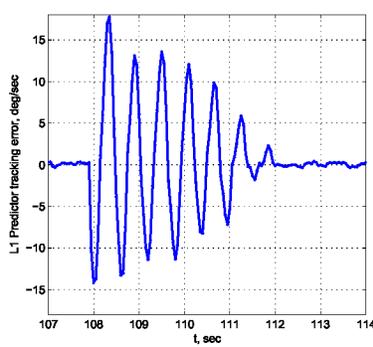


(c) \mathcal{L}_1 contribution

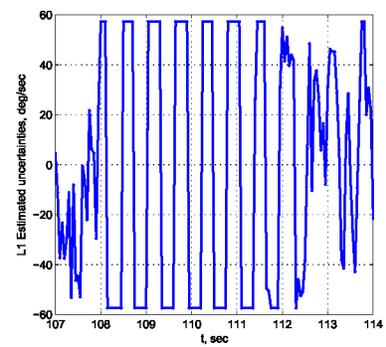
Figure 9: \mathcal{L}_1 . Closed-loop response in the presence of bending modes to biased sinusoidal reference signals.



(a) Output $r(t)$ vs Predictor Output $\hat{r}(t)$



(b) Tracking error $\tilde{r}(t)$



(c) uncertainty estimate $\hat{\sigma}(t)$

Figure 10: \mathcal{L}_1 . Saturation in the adaptive law.

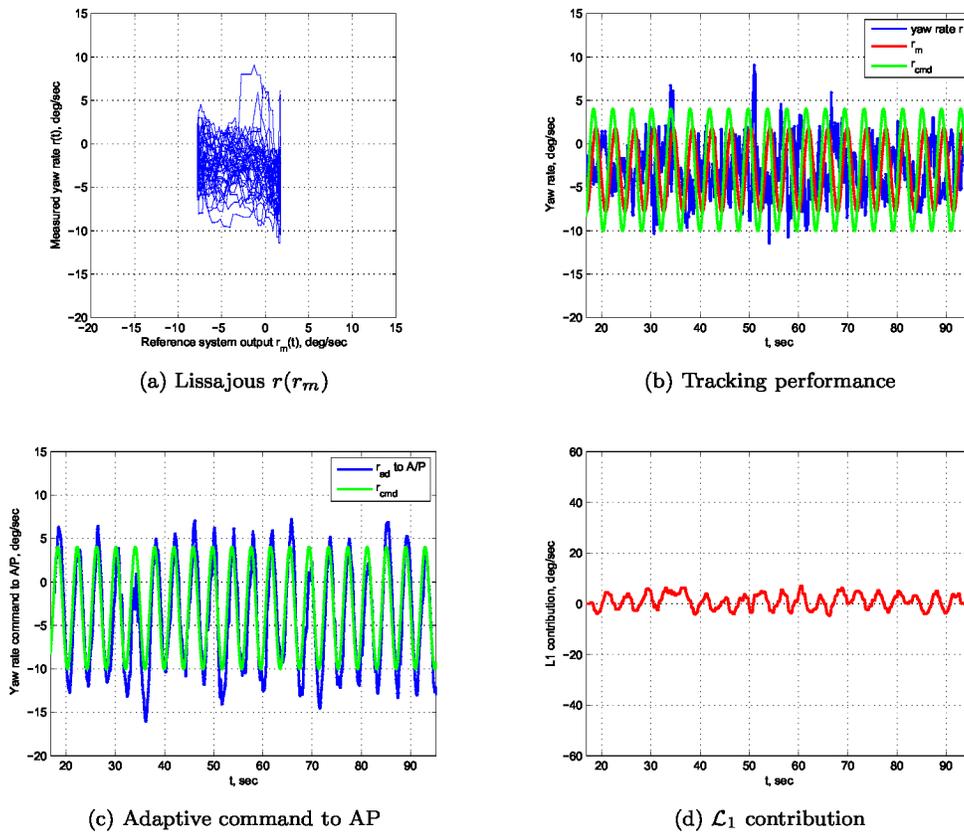


Figure 11: \mathcal{L}_1 . Closed-loop nominal response in the presence of second-order unmodeled dynamics to a biased sinusoidal reference signal at the phase crossover frequency.

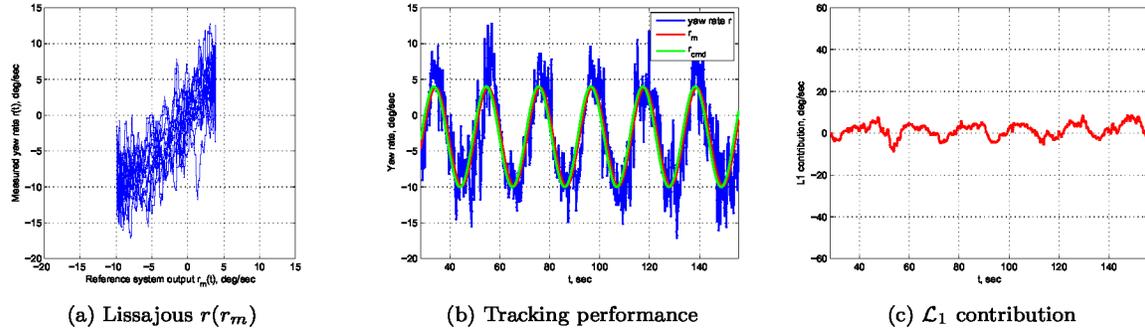


Figure 12: \mathcal{L}_1 . Closed-loop response in the presence of second-order unmodeled dynamics to biased sinusoidal reference signal at $\omega = 0.3 \frac{\text{rad}}{\text{s}}$.

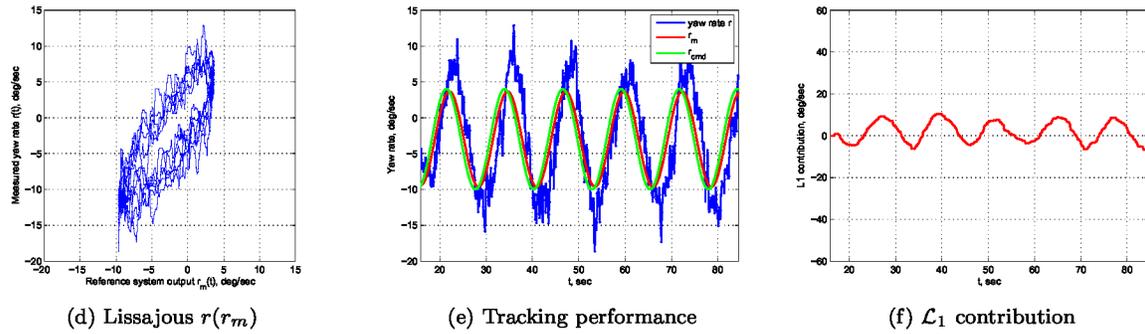


Figure 12: \mathcal{L}_1 . Closed-loop response in the presence of second-order unmodeled dynamics to biased sinusoidal reference signal at $\omega = 0.5 \frac{\text{rad}}{\text{s}}$.

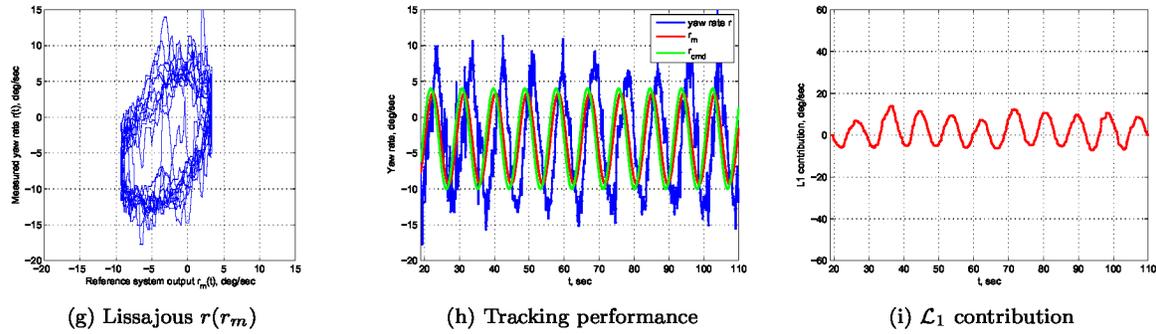


Figure 12: \mathcal{L}_1 . Closed-loop response in the presence of second-order unmodeled dynamics to biased sinusoidal reference signal at $\omega = 0.7 \frac{\text{rad}}{\text{s}}$.

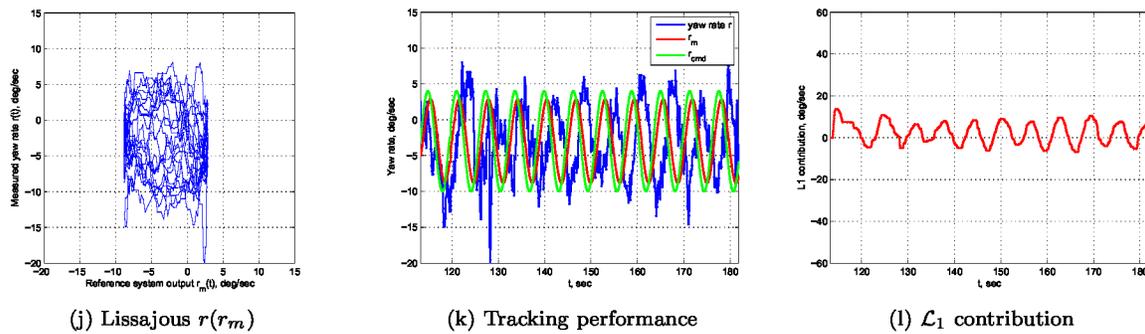


Figure 12: \mathcal{L}_1 . Closed-loop response in the presence of second-order unmodeled dynamics to biased sinusoidal reference signal at $\omega = 1.0 \frac{\text{rad}}{\text{s}}$.

References

- ¹Jacklin, S. A., Lowry, M. R., Schumann, J. M., Gupta, P. P., Bosworth, J. T., Zavala, E., Kelly, J. W., Hayhurst, K. J., Belcastro, C. M., and Belcastro, C. M., "Verification, Validation, and Certification Challenges for Adaptive Flight-Critical Control System Software," Providence, RI, August 2004, AIAA-2004-5258.
- ²Jacklin, S. A., "Closing Certification Gaps in Adaptive Flight Control Software," Honolulu, HI, August 2008, AIAA-2008-6988.
- ³Anderson, B. D. O. and Dehghani, A., "Challenges of Adaptive Control-past, permanent and future," *Annual Reviews in Control*, Vol. 32, No. 2, 2008, pp. 123–135.
- ⁴Wise, K. A., Lavretsky, E., Zimmerman, J., Francis, J., Dixon, D., and Whitehead, B., "Adaptive Flight Control of a Sensor Guided Munition," San Francisco, CA, August 2005, AIAA-2005-6385.
- ⁵Wise, K. A., Lavretsky, E., and Hovakimyan, N., "Adaptive Control in Flight: Theory, Application, and Open Problems," *American Control Conference*, Minneapolis, MN, June 2006, pp. 5966–5971.
- ⁶Cao, C., Patel, V. V., Reddy, C. K., Hovakimyan, N., Lavretsky, E., and Wise, K. A., "Are Phase and Time-delay Margins Always Adversely Affected by High-Gain?" Keystone, CO, August 2006, AIAA-2006-6347.
- ⁷Cao, C. and Hovakimyan, N., "Design and Analysis of a Novel \mathcal{L}_1 Adaptive Control Architecture with Guaranteed Transient Performance," *IEEE Transactions on Automatic Control*, Vol. 53, No. 3, 2008, pp. 586–591.
- ⁸Cao, C. and Hovakimyan, N., "Stability Margins of \mathcal{L}_1 Adaptive Control Architecture," *IEEE Transactions on Automatic Control*, Vol. 55, No. 2, February 2010.
- ⁹Cao, C. and Hovakimyan, N., " \mathcal{L}_1 Adaptive Controller for a Class of Systems with Unknown Nonlinearities: Part I," *American Control Conference*, Seattle, WA, June 2008, pp. 4093–4098.
- ¹⁰Zakrzewski, R. R., "Verification of Performance of a Neural Network Estimator," *IEEE World Congress on Computational Intelligence*, Honolulu, HI, May 2002.
- ¹¹Jordan, T. L., Langford, W. M., and Hill, J. S., "Airborne Subscale Transport Aircraft Research Testbed-Aircraft Model Development," San Francisco, CA, August 2005, AIAA-2005-6432.
- ¹²Xargay, E., Hovakimyan, N., and Cao, C., " \mathcal{L}_1 Adaptive Controller for Multi-Input Multi-Output Systems in the Presence of Nonlinear Unmatched Uncertainties," *American Control Conference*, Baltimore, MD, June–July 2010, Accepted for publication.
- ¹³Egardt, B., *Stability of Model Reference Adaptive and Self-Tuning Regulators*, Ph.D. thesis, Lund Institute of Technology, December 1978.
- ¹⁴Rohrs, C., Valavani, L., Athans, M., and Stein, G., "Robustness of continuous-time adaptive control algorithms in the presence of unmodeled dynamics," *Automatic Control, IEEE Transactions on*, Vol. 30, No. 9, 1985, pp. 881–889.
- ¹⁵Dobrokhodov, V., Kitsios, I., Kaminer, I., Jones, K., Xargay, E., Hovakimyan, N., Cao, C., Lizarraga, M., and Gregory, I., "Flight Validation of a Metrics Driven \mathcal{L}_1 Adaptive Control," *Proceedings of AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2008.